arXiv:1412.7922v1 [cs.DC] 26 Dec 2014

Fast Partial Distance Estimation and Applications

Boaz Patt-Shamir, Tel Aviv University
Christoph Lenzen, Max-Planck Institute for Informatics

December 30, 2014

Abstract

We study approximate distributed solutions to the weighted *all-pairs-shortest-paths* (APSP) problem in the CONGEST model. We obtain the following results.

1. A deterministic $(1+o(1))$-approximation to APSP in $\tilde{\mathcal{O}}(n)$ rounds. This improves over the best previously known algorithm, by both derandomizing it and by reducing the running time by a $\Theta(\log n)$ factor.

In many cases, routing schemes involve relabeling, i.e., assigning new names to nodes and require that these names are used in distance and routing queries. It is known that relabeling is necessary to achieve running times of $o(n/\log n)$. In the relabeling model, we obtain the following results.

2. A randomized $\mathcal{O}(k)$-approximation to APSP, for any integer $k>1$, running in $\tilde{\mathcal{O}}(n^{1/2+1/k}+D)$ rounds, where D is the hop diameter of the network. This algorithm simplifies the best previously known result and reduces its approximation ratio from $\mathcal{O}(k\log k)$ to $\mathcal{O}(k)$. Also, the new algorithm uses uses labels of asymptotically optimal size, namely $\mathcal{O}(\log n)$ bits.
3. A randomized $\mathcal{O}(k)$-approximation to APSP, for any integer $k>1$, running in time $\tilde{\mathcal{O}}((nD)^{1/2}\cdot n^{1/k}+D)$ and producing *compact routing tables* of size $\tilde{\mathcal{O}}(n^{1/k})$. The node lables consist of $\mathcal{O}(k\log n)$ bits. This improves on the approximation ratio of $\Theta(k^2)$ for tables of that size achieved by the best previously known algorithm, which terminates faster, in $\tilde{\mathcal{O}}(n^{1/2+1/k}+D)$ rounds.

1 Introduction

To allow a network to be useful, it must facilitate routing messages between nodes. By the very nature of networks, this computation must be distributed, i.e., there must be a distributed algorithm that computes the local data structures that support routing at network junctions (i.e., routing tables at nodes). A trivial distributed algorithm for this purpose is to collect the entire topology at a single location, apply a centralized algorithm, and distribute the result via the network. This simplistic approach is costly, in particular if the available bandwidth is limited. To study the distributed time complexity of routing table computation, we use the view point of the CONGEST model, i.e., we assume that each link in an n-node network allows only for $\mathcal{O}(\log n)$ bits to be exchanged in each unit of time.

In this work, we consider networks modeled by weighted undirected graphs, where the edge weights represent some abstract link cost, e.g., latency. Regarding routing, which is is a generic task with many variants, we focus on the following specific problems.

- *All-pairs distance estimation*: How fast can each node obtain an estimate of the distance to each other node, and how good is that estimate?
- *All-Pairs Shortest Paths:* How fast can we construct local data structures so that when given a destination node identifier, the node can locally determine the next hop on a path to the destination, and what's the *stretch* of the resulting route w.r.t. the shortest path?

The variants above are graph-theoretic; in modern routing systems, it is common practice to assign to nodes labels (identifiers) that contain some routing information. IP addresses, for example, contain a "network" part and a "host" part, which allow for hierarchical routing. Thus, the following additional questions are also of interest to us.

- *Routing Table Construction:* What are the answers to the above questions if we allow *relabeling*, i.e., allow the algorithm to choose (small) *labels* as node identifier, and require that distance and routing queries refer to nodes using these labels?
- *Compact Routing:* What are the answers to the above questions if we require that storage space allocated at the nodes (routing table size) is small?

Background. Shortest paths are a central object of study since the dawn of the computer era. The Bellman-Ford algorithm [4, 7], although originally developed for centralized optimization, is one of the few most fundamental distributed algorithms. Implemented as RIP, the algorithm was used in the early days of the Internet (when it was still called ARPANET) [12]. Measured in terms of the CONGEST model, a Bellman-Ford all-pairs shortest paths computation in weighted graphs takes $\Theta(n^2)$ time in the worst case, and requires $\Theta(n \log n)$ bits of storage at each node. Another simple solution to the problem is to collect the complete topology at each node (by flooding) and then apply a local single-source shortest paths algorithm, such as Dijkstra's. This solution has time complexity $\Theta(m)$ and storage complexity $\Theta(m)$, where m denotes the number of links in the network. It also enjoys improved stability and flexibility, and due to these reasons it became the Internet's routing algorithm in its later stages of evolution (see discussion in [11]). It is standardized as OSPF [13], which contains, in addition, provisions for hierarchical routing.

State of the art and New Results. Recently there has been a flurry of new results about routing in the CONGEST model. Instead of trying to track them all, let us start by reviewing known lower bounds, which help placing our results in the context of what is possible.

- Without relabeling, any polylogarithmic-ratio approximation to APSP requires $\tilde{\Omega}(n)$ rounds [14, 15].[1] This holds also if tables must only enable either distance estimates or routing, but not both.
- With node relabeling, any approximation to APSP requires $\tilde{\Omega}(\sqrt{n} + D)$ rounds [6], where D denotes the *hop diameter* of the network (see Section 2.2. The bound holds for both routing and distance

[1] Throughout this paper, we use $\tilde{\mathcal{O}}$-notation, which hides poly-logarithmic factors. See Section 2.6.

queries, and even for $D \in \mathcal{O}(\log n)$. (However, if routing may be *stateful*, i.e., routing decisions may depend on the tables of previously visited nodes, no non-trivial lower bound is known; all our routing algorithms are stateless.)

- If the routing table size is $\tilde{\mathcal{O}}(n^{1/k})$, then the approximation ratio of the induced routes is at least $2k-1$ [1, 18]. (This result does not hold for stateful routing.) For distance approximation, the same bound has been established for the special cases of $k \in \{1,2,3,5\}$, and is conjectured to hold for any k (see [21] and references therein).
- It is known [9] that any randomized $(2-o(1))$-approximation of APSP, and that any $(2-o(1))$-approximation of the weighted diameter of a graph takes $\tilde{\Omega}(n)$ time in the worst case. In the *unweighted* case, [8] show an $\tilde{\Omega}(n)$ lower bound on the time required to approximate the diameter to within a $3/2$ factor.

Let us now review the best known upper bounds and compare them with our results.

- For any $\varepsilon > 0$, we give a *deterministic* $(1+\varepsilon)$-approximation to APSP that runs in $\mathcal{O}(\varepsilon^{-2} n \log n)$ rounds. The best known previous result, due to Nanongkai [14], achieves the same approximation ratio within $\mathcal{O}(\varepsilon^{-2} n \log^2 n)$ rounds with high probability—Nanongkai's algorithm is randomized. We note that independently and concurrently to our work, Holzer and Pinsker [9] derived the same algorithm and result, and applied it in the Broadcast Congested Clique model, in which in each round, each node posts a single $\mathcal{O}(\log n)$-bit message which is delivered to all other nodes.
- For any $k \in \mathbb{N}$, we obtain a randomized $(6k-1+o(1))$-approximation to APSP running in time $\tilde{\mathcal{O}}(n^{1/2+1/(4k)} + D)$. The algorithm succeeds with high probability (see Section 2.6), as do all our randomized algorithms. This improves our previous work [15] by simplifying it and by reducing the approximation ratio from $\mathcal{O}(k \log k)$ to $\mathcal{O}(k)$. The new algorithm relabel nodes with labels of $\mathcal{O}(\log n)$ bits, whereas the previous one required $\mathcal{O}(\log n \log k)$-bit labels.
- For any $k \in \mathbb{N}$, we give a randomized $(4k-3+o(1))$-approximation to APSP running in time $\tilde{\mathcal{O}}(\min\{(nD)^{1/2} \cdot n^{1/k}, n^{2/3+2/(3k)} + D)$ with tables of size $\tilde{\mathcal{O}}(n^{1/k})$. This improves over the stretch of $\mathcal{O}(k^2)$ in own previous work [15], at the cost of increasing the running time (from $\tilde{\mathcal{O}}(n^{1/2+1/k} + D)$). We point out, however, that the proposed algorithm is the first that achieves an asymptotically optimal trade-off between stretch and table size in time $\tilde{o}(n)$ for all graphs of diameter $D \in \tilde{o}(n)$ and $k > 2$.

Technical Discussion. Our key tool is a generalization of the (S,h,σ)-*detection problem*, introduced in [10].[2] The problem is defined as follows. Given a graph with a distinguished set of *source nodes* S, the task is for each node to find the distances to its closest $\sigma \in \mathbb{N}$ sources within $h \in \mathbb{N}$ hops (the formal definition is given in Section 2.5). In [10] it is shown that this task can be solved in $h+\sigma$ rounds on *unweighted* graphs. The main new ingredient in all our results is an algorithm that, within a comparable running time, produces an *approximate* solution to (S,h,σ)-detection in *weighted* graphs.

Weighted graphs present a significant difficulty, because in weighted graphs, the number of hops in a shortest (by weight) path between two nodes may be a factor of $\Theta(n)$ larger than the minimal number of hops on any path connecting the same two nodes (the Congested Clique provides an extreme example of this phenomenon). Therefore, naively finding the absolute closest σ sources (w.r.t. weighted distance) within h hops may require running time $\Omega(n)$ in the worst case, for any h and σ. One may circumvent this difficulty by replacing the underlying graph metric by *h-hop distances*, which for $v,w \in V$ is defined as the minimum weight of all v-w-paths that consist of at most h hops. The collection of h-hop distances does not constitute a metric, but one can solve the (S,h,σ)-detection problem under h-hop distances in time σh using techniques similar to those used in the unweighted case [15].

Unfortunately, as illustrated in Figure 1, this time complexity is optimal in the worst case. To avoid this bottleneck, Nanongkai [14] used a rounding technique that had previously been leveraged in the centralized

[2] In the original paper, the third parameter is called k. We use σ here to avoid confusion with the use of k as the parameter controlling the trade-off between approximation ratio and table size.

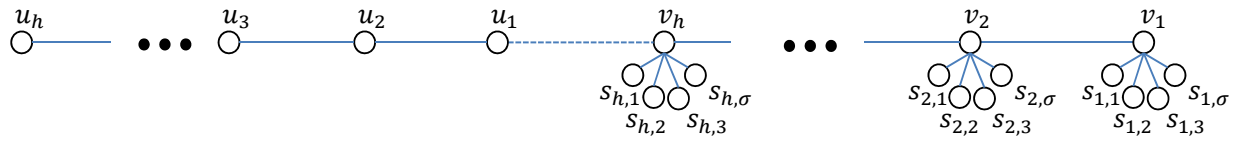

Figure 1: A graph where $(S, h+1, \sigma)$-detection cannot be solved in $o(h\sigma)$ rounds. Edge weights are $4ih$ for edges $\{v_i, s_{i,j}\}$ for all $i \in \{1, \ldots, h\}$ and $j \in \{1, \ldots, \sigma\}$, and 1 (i.e., negligible) for all other edges. Node u_i, $i \in \{1, \ldots, h\}$, needs to learn about all nodes $s_{i,j}$ and distances $\mathrm{wd}_{h+1}(u_i, s_{i,j})$, where $j \in \{1, \ldots, \sigma\}$. Hence all this information must traverse the dashed edge $\{u_1, v_h\}$. The example can be modified to attach the same source set to each v_h. Varying distances, then still $\sigma h = |S|h$ values must be communicated over the dashed edge. Hence, the special case $|S| = \sigma$ is not easier.

setting [22], solving the problem to within $(1+\varepsilon)$ factor by essentially reducing the weighted instance to $\mathcal{O}(\log n/\varepsilon)$ unweighted instances. In [14], the idea is to solve each instance using breadth-first-search. To avoid collisions, a random delay is applied to the starting time of each instance. The result is $(1+\varepsilon)$-approximate distances to all sources in $\mathcal{O}((h+|S|)\log^2 n/\varepsilon^2)$ rounds w.h.p. We replace this part of the algorithm with the deterministic source detection algorithm from [10], obtaining a deterministic algorithm of running time $\mathcal{O}((h+|S|)\log n)$. This, using $S = V$ and $h = \sigma = n$, has the immediate corollary of a deterministic $(1+o(1))$-approximation to APSP.

To derive our other results, solving the special case of $\sigma = |S|$ is insufficient. Consequently, we define a $(1+\varepsilon)$-approximate version of the (S, h, σ)-detection problem which we call *partial distance estimation*, or *PDE* for short (see Definition 2.2). The crucial insight is that by combining Nanongkai's and Zwick's rounding scheme with the algorithm from [10], PDE can be solved in $\mathcal{O}((h+\sigma)\log n/\varepsilon^2)$ rounds, with a bound of $\mathcal{O}(\sigma^2)$ on the number of messages sent by each node throughout the computation. Exploiting these properties carefully, we obtain our other results.

Further ingredients. Our compact routing schemes are distributed constructions of the routing hierarchies of Thorup and Zwick [20]. These make use of efficient tree labeling schemes presented in the same paper, which allow for a distributed implementation in time $\tilde{\mathcal{O}}(h)$ in trees of depth h if relabeling is permitted. For compact routing table construction, we continue the Thorup-Zwick construction by simulating the partial distance estimation algorithm on the *skeleton graph* [15], broadcasting all messages via a BFS tree. This avoids the quadratic stretch incurred by the approach in [15] due to approximating distances in the skeleton graph using a *spanner* [17], which is constructed by simulating the Baswana-Sen algorithm [3]. If compact tables are not required, the partial distance estimation algorithm enables to collapse the Thorup-Zwick hierarchy of the lower levels into a single step, giving rise to a constant approximation ratio, and thus removing the multiplicative overhead of $\mathcal{O}(\log k)$ from [15] in constructing non-compact routing tables.

Organization of the paper. In Section 2 we define the basic graph-theoretic concepts, problems, and execution model. In Section 3 we describe the algorithm for partial distance estimation. In Section 4 we present the improved results we derive using the algorithm for PDE.

2 Model and Problem

2.1 Execution Model

We follow the CONGEST model as described by [16]. The distributed system is represented by a simple, connected weighted graph $G = (V, E, W)$, where V is the set of nodes, E is the set of edges, and $W : E \to \mathbb{N}$ is the edge weight function. As a convention, we use n to denote the number of nodes. We assume that all edge weights are bounded by some polynomial in n, and that each node $v \in V$ has a unique identifier of

$\mathcal{O}(\log n)$ bits (we use v to denote both the node and its identifier).

Execution proceeds in global synchronous rounds, where in each round, each node takes the following three steps:

1. Perform local computation,
2. send messages to neighbors, and
3. receive the messages sent by neighbors.

Initially, nodes are aware only of their neighbors; input values (if any) are assumed to be fed by the environment before the first round. Throughout this paper, we assume that node v is given the weight of each edge $\{v, w\} \in E$ as input. Output values, which are computed at the end of the final round, are placed in special output-registers. In each round, each edge can carry a message of B bits for some given parameter B of the model; we assume that $B \in \Theta(\log n)$ throughout this paper.

2.2 Graph-Theoretic Concepts

Fix a weighted undirected graph $G = (V, E, W)$. A *path* p connecting $v, w \in V$ is a sequence of nodes $\langle v = v_0, \ldots, v_k = w \rangle$ such that for all $0 \leq i < k$, $\{v_i, v_{i+1}\}$ is an edge in G. Let $\mathrm{paths}(v, w)$ denote the set of all paths connecting nodes v and w. We use the following unweighted concepts.

- The *hop-length* of a path p, denoted $\ell(p)$, is the number of edges in it.
- A path p_0 between v and w is a *shortest unweighted path* if its hop-length $\ell(p_0)$ is minimum among all $p \in \mathrm{paths}(v, w)$.
- The *hop distance* $\mathrm{hd} : V \times V \to \mathbb{N}_0$ is defined as the hop-length of a shortest unweighted path, $\mathrm{hd}(v, w) := \min\{\ell(p) \mid p \in \mathrm{paths}(v, w)\}$.
- The *hop-diameter* of a graph $G = (V, E, W)$ is $\mathrm{D} \stackrel{\mathrm{def}}{=} \max_{v,w \in V}\{\mathrm{hd}(v, w)\}$.

We use the following weighted concepts.

- The *weight* of a path p, denoted $W(p)$, is its total edge weight, i.e., $W(p) \stackrel{\mathrm{def}}{=} \sum_{i=1}^{\ell(p)} W(v_{i-1}, v_i)$.
- A path p_0 between v and u is a *shortest weighted path* if its weight $W(p_0)$ is minimum among all $p \in \mathrm{paths}(v, w)$.
- The *weighted distance* $\mathrm{wd} : V \times V \to \mathbb{N}$ is defined as the weight of a shortest weighted path, $\mathrm{wd}(v, u) \stackrel{\mathrm{def}}{=} \min\{W(p) \mid p \in \mathrm{paths}(v, u)\}$.
- The *weighted diameter* of G is $\mathrm{WD} \stackrel{\mathrm{def}}{=} \max\{\mathrm{wd}(v, u) \mid v, u \in V\}$.

Finally, we define the notion of the *shortest paths distance*.

- A path p_0 between v and w is a *minimum-hop shortest weighted path* if its hop-length $\ell(p_0)$ is minimum among all shortest weighted paths connecting v and w. The number of hops in such a path $h_{v,w}$ is the *shortest path distance* of v and w.
- The *shortest path diameter* of G is $\mathrm{SPD} \stackrel{\mathrm{def}}{=} \max\{h_{v,w} \mid v, w \in V\}$.

2.3 Routing

In the *routing table construction* problem (abbreviated RTC), the output at each node v consists of (i) a unique *label* $\lambda(v)$ and (ii) a function "next_v" that takes a destination label λ and produces a neighbor of v, such that given the label $\lambda(w)$ of any node w, and starting from any node v, it is possible to reach w from v by following the next pointers. Formally, the requirement is as follows. Given a start node v and a destination label $\lambda(w)$, let $v_0 = v$ and define $v_{i+1} = \mathrm{next}_{v_i}(\lambda(w))$ for $i \geq 0$. Then $v_i = w$ for some i.

The performance of a solution is measured in terms of its *stretch*. A route is said to have stretch $\rho \geq 1$ if its total weight is ρ times the weighted distance between its endpoints, and a solution to RTC is said to have stretch ρ if the maximal stretch of all its induced routes is ρ.

Variants. Routing appears in many incarnations. We list a few important variants below.

Name-independent routing. Our definition of RTC allows for node relabeling. This is the case, as mentioned above, in the Internet. The case where no such relabeling is allowed (which can be formalized by requiring λ to be the identity function), was introduced in [2] as *name-independent* routing. Note that any scheme can be trivially transformed into a name-independent one by announcing all node/label pairs, but this naive approach requires to broadcast and store $\Omega(n \log n)$ bits.

Stateful routing. The routing problem as defined above is *stateless* in the sense that routing a packet is done regardless of the path it traversed so far. One may also consider *stateful* routing, where while being routed, a packet may gather information that helps it navigate later. (One embodiment of this idea in the Internet routing today is MPLS, where packets are temporarily piggybacked with extra headers. Early solutions to the compact routing problem, such as that of [19], were also based on stateful routing.) Note that the set of routes to a single destination in stateless routing must constitute a tree, whereas in stateful routing even a single route may contain a cycle. Formally, in stateful routing the label of the destination may change from one node to another: The next_v function outputs both the next hop (a neighbor node), and a new label λ_v used in the next hop.

2.4 Distance Approximation

The *distance approximation* problem is closely related to the routing problem. Again, each node v outputs a label $\lambda(v)$, but now, v needs to construct a function $\mathrm{dist}_v : \lambda(V) \to \mathbb{R}^+$ (the table) such that for all $w \in V$ it holds that $\mathrm{dist}_v(\lambda(w)) \geq \mathrm{wd}(v,w)$. The stretch of the approximation v has of $\mathrm{wd}(v,w)$ is $\mathrm{dist}_v(\lambda(w))/\mathrm{wd}(v,w)$, and the solution has stretch $\rho \geq 1$ if $\max_{v,w \in V}\{\mathrm{dist}_v(\lambda(w))/\mathrm{wd}(v,w)\} = \rho$.

Similarly to routing, we call a scheme name-independent if λ is the identity function. Since we require distance estimates to be produced without communication, there is no "stateful" distance approximation.

2.5 Partial Distance Estimation (PDE)

The basic problem we attack in this paper is partial distance estimation, which generalizes the source detection problem. Let us start by defining the simpler variant.

Given a set of nodes $S \subseteq V$ and a parameter $h \in \mathbb{N}$, $L_v^{(h)}$ denotes the list resulting from ordering the set $\{(\mathrm{wd}(v,w),w) \mid w \in S \wedge h_{v,w} \leq h\}$ lexicographically in ascending order, i.e.,

$$(\mathrm{wd}(v,w),w) < (\mathrm{wd}(v,u),u) \Leftrightarrow \big((\mathrm{wd}(v,w) < \mathrm{wd}(v,u)) \vee (\mathrm{wd}(v,w) = \mathrm{wd}(v,u) \wedge w < u)\big).$$

Definition 2.1 ((S,h,σ)-detection)**.** *Given are a set of* sources $S \subseteq V$ *and parameters* $h, \sigma \in \mathbb{N}$*. Each node is assumed to know* h*,* σ*, and whether it is in* S *or not. The goal is to compute at each node* $v \in V$ *the list* L_v *of the top* σ *entries in* $L_v^{(h)}$*, or the complete* $L_v^{(h)}$ *if* $|L_v^{(h)}| \leq \sigma$*.*

Relaxing this by allowing approximation to within $(1+\varepsilon)$, we arrive at the following definition.

Definition 2.2 (Partial Distance Estimation (PDE))**.** *Given* $S \subseteq V$*,* $h, \sigma \in \mathbb{N}$*, and* $\varepsilon > 0$*,* $(1+\varepsilon)$-approximate (S,h,σ)-estimation *is defined as follows. Determine a* distance function $\mathrm{wd}' : V \times S \to \mathbb{N} \cup \infty$ *satisfying*

- $\forall v \in V, s \in S:\ \mathrm{wd}'(v,s) \geq \mathrm{wd}(v,s)$*, where* wd *is the weighted distance from* v *to* s*, and*
- *if* $h_{v,s} \leq h$*, then* $\mathrm{wd}'(v,s) \leq (1+\varepsilon)\,\mathrm{wd}(v,s)$*.*

For each $v \in V$*, sort the set* $\{(\mathrm{wd}'(v,s),s) \mid s \in S\}$ *in ascending lexicographical order. Each node* v *needs to output the prefix* L_v *of the sorted list consisting of the first (up to)* σ *elements with* $\mathrm{wd}'(v,s) < \infty$*.*

This generalizes the source detection problem in that setting $\varepsilon = 0$ and choosing wd' as h-hop distances results in an exact weighted version of the source detection problem, and, specializing further to unweighted graphs, h-hop distances just become hop distances to nodes within h hops.

2.6 General Concepts

We use extensively "soft" asymptotic notation that ignores polylogarithmic factors. Formally, $g(n) \in \tilde{\mathcal{O}}(f(n))$ if and only if there exists a constant $c \in \mathbb{R}_0^+$ such that $f(n) \leq g(n) \log^c n$ for all but finitely many values of $n \in \mathbb{N}$. Analogously,

- $f(n) \in \tilde{\Omega}(g(n))$ if and only if $g(n) \in \tilde{\mathcal{O}}(f(n))$,
- $\tilde{\Theta}(f(n)) \stackrel{\text{def}}{=} \tilde{\mathcal{O}}(f(n)) \cap \tilde{\Omega}(f(n))$,
- $g(n) \in \tilde{o}(f(n))$ if and only if for each $c \in \mathbb{R}_0^+$ it holds that $\lim_{n\to\infty} g(n) \log^c(f(n))/f(n) = 0$, and
- $g(n) \in \tilde{\omega}(f(n))$ if and only if $f(n) \in \tilde{o}(g(n))$.

To model probabilistic computation, we assume that each node has access to an infinite string of independent unbiased random bits. When we say that a certain event occurs "with high probability" (abbreviated "w.h.p."), we mean that the probability of the event not occurring can be set to be less than $1/n^c$ for any desired constant c, where the probability is taken over the strings of random bits. Due to the union bound, this definition entails that any polynomial number of events that occur w.h.p. also jointly occur w.h.p. We will make frequent use of this fact throughout the paper.

3 From Weighted to Unweighted

Fix $0 < \varepsilon \in \mathcal{O}(1)$. Using the technique of Nanongkai [14], we reduce PDE to $\mathcal{O}(\log_{1+\varepsilon} \mathrm{WD})$ instances of the unweighted problem as follows. Let $i_{\max} \stackrel{\text{def}}{=} \log_{1+\varepsilon} w_{\max}$, where $w_{\max}$ is the largest edge weight in G. Note that since we assume that edge weights are polynomial in n, $i_{\max} \in \mathcal{O}(\log_{1+\varepsilon} n)$. Clearly $i_{\max}$ can be determined in $\mathcal{O}(D)$ rounds.

For $i \in \{1, \ldots, i_{\max}\}$, let $b(i) \stackrel{\text{def}}{=} (1+\varepsilon)^i$, and define $W_i : E \to b(i) \cdot \mathbb{N}$ by $W_i(e) \stackrel{\text{def}}{=} b(i)\lceil W(e)/b(i) \rceil$, i.e., by rounding up edge weights to integer multiples of $(1+\varepsilon)^i$. Denote by wd_i the resulting distance function, i.e., the distance function of the graph (V, E, W_i). Then the following crucial property holds.

Lemma 3.1 (adapted from [14])**.** *For all $v, w \in V$ and*

$$i_{v,w} := \max\left\{0, \left\lfloor \log_{1+\varepsilon}\left(\frac{\varepsilon \,\mathrm{wd}(v,w)}{h_{v,w}}\right)\right\rfloor\right\},$$

it holds that

$$\mathrm{wd}_{i_{v,w}}(v,w) < (1+\varepsilon)\,\mathrm{wd}(v,w) \in \mathcal{O}\left(\frac{b(i_{v,w})h_{v,w}}{\varepsilon}\right).$$

Proof. W.l.o.g., $i_{v,w} \neq 0$, as $b_0 = 1$ and hence $\mathrm{wd}_0 = \mathrm{wd}$. The choice of $i_{v,w}$ yields that

$$\mathrm{wd}_{i_{v,w}}(v,w) < \mathrm{wd}(v,w) + b(i_{v,w})h_{v,w} \leq (1+\varepsilon)\,\mathrm{wd}(v,w).$$

To see the second bound, note that by definition of $i_{v,w}$ and $b(i_{v,w})$,

$$\mathrm{wd}(v,w) \leq \frac{(1+\varepsilon)b(i_{v,w})h_{v,w}}{\varepsilon}.$$

Due to the previous inequality and the constraint that $\varepsilon \in \mathcal{O}(1)$, the claim follows. □

Next, let G_i be the *unweighted* graph obtained by replacing each edge e in (V, E, W_i) by a path of $W_i(e)/b(i)$ edges. Let $\mathrm{hd}_i(v,w)$ denote the distance (minimal number of hops) between v and w in G_i. The previous lemma implies that in $G_{i_{v,w}}$, the resulting hop distance between v and w is not too large.

Corollary 3.2. *For each $v, w \in V$, it holds that* $\mathrm{hd}_{i_{v,w}}(v,w) \in \mathcal{O}(h_{v,w}/\varepsilon)$.

Proof. By Lemma 3.1, $\mathrm{wd}_{i_{v,w}}(v,w) \in \mathcal{O}(b(i_{v,w})h_{v,w}/\varepsilon)$. As edge weights are scaled down by factor $b(i_{v,w})$ in $G_{i_{v,w}}$, this implies $\mathrm{hd}_{i_{v,w}}(v,w) \in \mathcal{O}(h_{v,w}/\varepsilon)$. □

These simple observations imply that an efficient algorithm for unweighted source detection can be used to solve partial distance estimation at the cost of a small increase in running time.

Theorem 3.3. *Any deterministic algorithm for unweighted (S,h,σ)-detection with running time $R(h,\sigma)$ can be employed to solve $(1+\varepsilon)$-approximate (S,h,σ)-estimation in $\mathcal{O}(\log_{1+\varepsilon} n \cdot R(h',\sigma) + D)$ rounds, for some $h' \in \mathcal{O}(h/\varepsilon)$.*

Proof. Let $\mathcal{A}$ be any deterministic algorithm for unweighted (S,h,σ)-detection with running time $R(h,\sigma)$. We use the following algorithm for partial distance estimation.

1. Let h' be such that $\mathrm{hd}_{i_{v,s}}(v,s) \leq h'$ for all $v \in V$ and $s \in S$ with $h_{v,s} \leq h$. By Corollary 3.2, there is such an $h' \in \mathcal{O}(h/\varepsilon)$.
2. For $i \in \{1,\dots,i_{\max}\}$, solve (S,h',σ)-detection on G_i by $\mathcal{A}$. Denote by $L_{v,i}$ the computed list.
3. For $s \in S$, define

$$\tilde{\mathrm{wd}}(v,s) \stackrel{\text{def}}{=} \inf\{\mathrm{hd}_i(v,s)b(i) \mid 0 \leq i \leq i_{\max} \wedge (\mathrm{hd}_i(v,s),s) \in L_{v,i}\}.$$

   Note that if $\tilde{\mathrm{wd}}(v,s) < \infty$, then v can determine s and $\tilde{\mathrm{wd}}(v,s)$ from the previous step. Each node v outputs the list L_v consisting of the (up to) first σ elements of the set $\{(\tilde{\mathrm{wd}}(v,s),s) \mid \tilde{\mathrm{wd}}(v,s) < \infty\}$, with respect to ascending lexicographical order.

Clearly, the resulting running time is the one stated in the claim of the theorem.

In order to show that the output is feasible, i.e., satisfies the guarantees of partial distance estimation with approximation ratio $1+\varepsilon$, define

$$\mathrm{wd}'(v,s) \stackrel{\text{def}}{=} \inf\{\mathrm{hd}_i(v,s)b(i) \mid 0 \leq i \leq i_{\max} \wedge \mathrm{hd}_i(v,s) \leq h'\}.$$

We claim that the required properties are satisfied with respect to wd' and that the list returned by v is the one induced by wd', which will complete the proof. The claim readily follows from the following properties.

1. $\forall v \in V, s \in S:\ \mathrm{wd}'(v,s) \geq \mathrm{wd}(v,s)$,
2. $\forall v \in V, s \in S:\ h_{v,s} \leq h \Rightarrow \mathrm{wd}'(v,s) \leq (1+\varepsilon)\,\mathrm{wd}(v,s)$,
3. $\forall v \in V, s \in S:\ (\tilde{\mathrm{wd}}(v,s),s) \geq (\mathrm{wd}'(v,s),s)$, and
4. $\forall v \in V, (\tilde{\mathrm{wd}}(v,s),s) \in L_v:\ \tilde{\mathrm{wd}}(v,s) = \mathrm{wd}'(v,s)$.

Hence, it remains to show these four properties.

1. By definition, $b(i)\,\mathrm{hd}_i(v,s) = \mathrm{wd}_i(v,s) \geq \mathrm{wd}(v,s)$ for all $v \in V$ and $s \in S$.
2. By the first step, $h_{v,s} \leq h \Rightarrow h_{i_{v,s}}(v,s) \leq h'$. Hence, $\mathrm{wd}'(v,s) \leq b(i_{v,s})h_{i_{v,s}}(v,s) = \mathrm{wd}_{i_{v,s}}(v,s) < (1+\varepsilon)\,\mathrm{wd}(v,s)$ by Lemma 3.1.
3. This trivially holds, because $(\mathrm{hd}_i(v,s),s) \in L_{v,i}$ implies that $\mathrm{hd}_i(v,s) \leq h'$ (we executed (S,h',σ)-detection on each G_i), i.e., $\tilde{\mathrm{wd}}(v,s)$ is an infimum taken over a subset of the set used for $\mathrm{wd}'(v,s)$.
4. Assume for contradiction that $(\tilde{\mathrm{wd}}(v,s),s) \in L_v$, yet $\tilde{\mathrm{wd}}(v,s) > \mathrm{wd}'(v,s)$ (by the previous property $\tilde{\mathrm{wd}}(v,s) < \mathrm{wd}'(v,s)$ is not possible). Choose i such that $b(i)\,\mathrm{hd}_i(v,s) = \mathrm{wd}'(v,s)$ and $\mathrm{hd}_i(v,s) \leq h'$. We have that $(\mathrm{hd}_i(v,s),s) \notin L_{v,i_{v,s}}$, as otherwise we had $\tilde{\mathrm{wd}}(v,s) \leq b(i)\,\mathrm{hd}_i(v,s) = \mathrm{wd}'(v,s)$. It follows that $|L_{v,i}| = \sigma$ and, for each $(\mathrm{hd}_i(v,t),t) \in L_{v,i}$, we have that

$$(\tilde{\mathrm{wd}}(v,t),t) \leq (b(i)\,\mathrm{hd}_i(v,t),t) < (b(i)\,\mathrm{hd}_i(v,s),s) = (\mathrm{wd}'(v,s),s) \leq (\tilde{\mathrm{wd}}(v,s),s),$$

   where in the final step we exploit the third property. As there are σ distinct such sources σ, we arrive at the contradiction that $(\tilde{\mathrm{wd}}(v,s),s) \notin L_v$. □

Lemma 3.4. *Consider the unweighted source detection algorith of [10]. The number of messages a node v broadcasts (i.e., sends to all neighbors) until it announces the first up to σ elements from $L_v^{(h)}$ is at most $\mathcal{O}(\sigma^2)$. Moreover, the solution to PDE given in Theorem 3.3 can be implemented so that each node sends $\tilde{\mathcal{O}}(\sigma^2)$ messages.*

Proof. Suppose the i^{th} element of L_v is $(d(v,s),s)$. By Lemma 4.2 from [10], v does not send $(d(v,s),s)$ in any round $r > d(v,s)+i$.[3] As it does send this value eventually, it must do so at the latest in round $d(v,s)+i$. Afterwards, no further messages (d_s,s) will be sent by v, as only smaller distances could be announced, but the minimum value $d(v,s)$ was already transmitted. Moreover, v cannot send any message (d_s,s) earlier than round $d(v,s)+1$, as $d(v,s)$ rounds are required for v to learn about the existence of s. Hence, the total number of messages concerning the first (up to) σ elements of $L_v^{(h)}$ during the course of the algorithm are at most $\sum_{i=1}^{\sigma} i \in \Theta(\sigma^2)$. Denoting by h_v the distance of the σ^{th} element of $L_v^{(h)}$ (or h, if $|L_v^{(h)}| < \sigma$), all these messages have been sent by the end of round $h_v+\sigma$. Node v does not learn about sources in distance h_v or larger earlier than round h_v, hence it will have sent at most $\Theta(\sigma^2)$ messages by the end of round $h_v+\sigma$. As distances to other than the σ closest sources are also irrelevant to other nodes, the algorithm will still yield correct results if v stops sending any messages after this number of messages have been sent.

Using this slightly modified algorithm in the construction from Theorem 3.3, the second statement of the lemma follows. □

Corollary 3.5. *For any $0 < \varepsilon \in \mathcal{O}(1)$, $(1+\varepsilon)$-approximate (S,h,σ)-estimation can be solved in $\mathcal{O}((h+\sigma)/\varepsilon^2 \cdot \log n + D)$ rounds. Tables of size $\mathcal{O}(\sigma \log n)$ for routing with stretch $1+\varepsilon$ from each $v \in V$ to the (up to) σ detected nodes can be constructed in the same time. Each node broadcasts (i.e., sends the same message to all neighbors) in $\mathcal{O}(\sigma^2/\varepsilon \cdot \log n)$ rounds during the course of the algorithm.*

Proof. We apply Theorem 3.3 to the algorithm from [10], which has running time $R(h',\sigma) = h' + \sigma \in \mathcal{O}(h/\varepsilon + \sigma)$. Since $\varepsilon \in \mathcal{O}(1)$, we have that $\mathcal{O}(\log_{1+\varepsilon} n) = \mathcal{O}(\log n/\varepsilon)$; this shows the first part of the claim. For the second, observe that if all nodes store their lists $L_{v,i}$ and send them to their neighbors, it is trivial to derive the respective routing tables. The third statement readily follows from Lemma 3.4. □

4 Applications

In this section we apply Corollary 3.5 to various aspects of routing in weighted graphs. We improve on the best known results for three questions: the running time required to compute small-stretch routes with and without node relabeling, and the stretch we can achieve within a given running time bound and routing table size.

4.1 Almost Exact APSP: Routing Without Node Relabeling

First we state our results for distributed computation of all-pairs $(1+\varepsilon)$-approximate shortest paths. The result follows simply by instantiating Corollary 3.5 with all nodes as sources and $h = \sigma = n$. As $h_{v,w} < n$ for all $v,w \in V$, $\mathrm{wd}'(v,w) \leq (1+\varepsilon)\,\mathrm{wd}(v,w) < \infty$. The returned lists thus contain entries for all $n = \sigma$ nodes.

Theorem 4.1. *$(1+\varepsilon)$-approximate APSP can be solved deterministically in $\mathcal{O}(n/\varepsilon^2 \cdot \log n)$ rounds.*

We note that Theorem 4.1 improves on the best known result for computing approximate shortest paths in the CONGEST model [14] in two ways: first, it is deterministic, and second, the running time is reduced by a logarithmic factor.

[3] There is a typographical error in the statement of the lemma in the paper; the inequality should read $d_s + \ell_v^{(r)}(d_s,s) < r$.

4.2 Routing Table Computation With Node Relabeling

In this section we use Corollary 3.5 to improve upon the best known previous result to compute routing tables when node relabeling is allowed [15]. We comment that node relabeling is a common practice in routing schemes: the idea is to encode some location information in the name of the nodes so as to reduce the space consumed by routing tables. For example, the Internet's IP addresses consist of a "network" part and a "host" part, which allows for hierarchical routing. In this case the label length is another performance measure that should be noted.

In [15], the algorithm guarantees the following. Given an integer $0 < k \leq \log n$, the algorithm computes (w.h.p.) in $\tilde{\mathcal{O}}(n^{1/2\cdot(1+1/k)}+D)$ rounds node labels of size $\mathcal{O}(\log n \log k)$ and routes with stretch $\mathcal{O}(k \log k)$. We will show how to execute this task without changing the running time, but improve the stretch and node label size by a $\log k$ factor.

First, let us briefly review the approach from [15].

1. Sample $\tilde{\Theta}(\sqrt{n})$ nodes uniformly, forming the *skeleton* S.
2. Construct and make known to all nodes an α-spanner of the *skeleton graph* (S, E_S, W_S). Here, $\{s,t\} \in E_S$ if $\mathrm{hd}(s,t) \leq h \in \tilde{\Theta}(\sqrt{n})$ and $W_S(s,t)$ is the minimum weight of an s-t path of at most h hops. It is shown that, w.h.p., distances in the skeleton graph are identical to distances in the original graph.
3. For each $v \in V$, denote by s_v the node in S minimizing $(\mathrm{wd}(v,s_v), s_v)$. For each node v, compute distances and routing tables with stretch β to all nodes $w \in V$ with $(\mathrm{wd}(v,w),w) \leq (\mathrm{wd}(v,s_v),s_v)$ and from s_v to v (using "tree routing"). This part is called the *short range* part of the scheme.
4. It is then shown that if $(\mathrm{wd}(v,w),w) > (\mathrm{wd}(v,s_v),s_v)$, routing on on a shortest paths from v to s_v, s_v to s_w in the skeleton spanner, (whose edges map to paths of the same weight in G), and finally from s_w to w has stretch $\mathcal{O}(\alpha\beta)$ (the same holds for distance estimation). This part is called *long range routing*.

To facilitate routing, the label of each node v contains the following components: the identity of the closest skeleton node[4] s_v and the distance to it $\mathrm{wd}(v,s_v)$; and a label for the short-range routing (required for the tree routing).

The spanner construction required for Step 2 can be used as black box, giving stretch $\alpha \in \Theta(1/k)$ within $\tilde{\mathcal{O}}(n^{1/2+1/k}+D)$ time. Moreover, it is known how to construct labels for tree routing of size $(1+o(1))\log n$ in time $\tilde{\mathcal{O}(h)}$ in trees of depth h [20].

We follow the general structure of [15] by implementing Step 3 so that $\beta \in \mathcal{O}(1)$ and the shortest-path trees do not become to deep. To this end, we apply Corollary 3.5 with $h = \sigma \approx \sqrt{n}$ and source set V. Note that the approximation error is not limited to giving approximate distances to the σ closest nodes: we may obtain distance estimates to an *entirely different* set of nodes. However, we can show that the distances of the nodes showing up in the list are at most factor $1+\varepsilon$ larger than their true distances. Denoting by $s'_v \in S$ denote the node minimizing $(\mathrm{wd}'(v,s),s)$ among nodes in S, the crucial properties we use are captured in the following lemma.

Lemma 4.2. *Suppose we sample each node into S with independent probability p and solve $(1+\varepsilon)$-approximate (V,h,σ)-estimation with $\min\{h,\sigma\} \geq c\log n/p$, where c is a sufficiently large constant. Then for all $v,w \in V$, w.h.p. the following statements hold.*

1. $(\mathrm{wd}'(v,w),w) \leq (\mathrm{wd}'(v,s'_v),s'_v) \Rightarrow \mathrm{wd}'(v,w) \leq (1+\varepsilon)\,\mathrm{wd}(v,w) \wedge (\mathrm{wd}(v,w),w) \in L_v$
2. $(\mathrm{wd}(v,w),w) \leq (\mathrm{wd}(v,s_v),s_v) \Rightarrow \mathrm{wd}'(v,w) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$
3. $(\mathrm{wd}(v,w),w) > (\mathrm{wd}(v,s_v),s_v) \Rightarrow \mathrm{wd}'(v,s_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$
4. $(\mathrm{wd}'(v,w),w) > (\mathrm{wd}'(v,s'_v),s'_v) \Rightarrow \mathrm{wd}'(v,s'_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$

[4] For convenience, we assume that always $S \neq \emptyset$, which holds w.h.p.

Proof. Fix $v \in V$ and order $\{(\mathrm{wd}'(v,w),w) \mid w \in V\}$ in ascending lexicographic order. Suppose $s'_v \in S$ is the i^{th} element of the resulting list. Note that, by definition, s'_v minimizes $(\mathrm{wd}'(v,s),s)$ among nodes $s \in S$. Hence, the probability that $i \geq \min\{h,\sigma\}$ is

$$(1-p)^{\min\{h,\sigma\}} \in e^{-\Theta(c\log n)} = n^{-\Theta(c)}.$$

Since c is a sufficiently large constant, this implies that $i < h$ w.h.p., and thus $h_{v,w} < i < \min\{h,\sigma\}$ for all w with $(\mathrm{wd}'(v,w),w) \leq (\mathrm{wd}'(v,s'_v),s'_v)$. By the properties of (V,h,σ)-estimation, it follows that, w.h.p., $\mathrm{wd}'(v,w) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$ and $(\mathrm{wd}'(v,w),w) \in L_v$ for all such w.

To show the second statement, we perform the same calculation for the list $\{(\mathrm{wd}(v,w),w) \mid w \in V\}$; the element from S minimizing $(\mathrm{wd}(v,s),s)$ is s_v. For the third statement, we apply the second to s_v, yielding that

$$\mathrm{wd}'(v,s_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,s_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$$

w.h.p. For the final statement, if $\mathrm{wd}'(v,w) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$, it follows that

$$\mathrm{wd}'(v,s'_v) \leq \mathrm{wd}'(v,w) \leq (1+\varepsilon)\,\mathrm{wd}(v,w).$$

Otherwise, the second statement shows that $\mathrm{wd}(v,w) \geq \mathrm{wd}(v,s_v)$ w.h.p., implying

$$\mathrm{wd}'(v,s'_v) \leq \mathrm{wd}'(v,s_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,s_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,w). \qquad \square$$

Our goal now is to let s'_v take the place of s_v in the original scheme. By the previous lemma and the time bound of $\tilde{\mathcal{O}}((h+\sigma)/\varepsilon^2)$ for $(1+\varepsilon)$-appromate (S,h,σ)-estimation, this achieves $\beta \in 1+o(1)$ within the desired time bound for $p \approx 1/\sqrt{n}$. However, this comes with a twist: as s'_v may be different from s_v, we must show that the resulting approximation ratio is still $\mathcal{O}(\alpha)$, and as we do not use exact distances, the $|S|$ trees induced by the approximately shortest paths from each v to s'_v might overlap. The following two lemmas address these issues, as well as the depth of the trees.

Lemma 4.3. *Suppose we sample each node into S with independent probability p and solve $(1+\varepsilon)$-approximate (S,h,σ)-estimation with $h = c\log n/p$, where c is a sufficiently large constant. Denote by wd'_S the respective distance function, and by wd' and L_v the distance function and output of $v \in V$, respectively, of a solution to $(1+\varepsilon)$-approximate (V,h,h)-estimation. If for $v,w \in V$ it holds that $(\mathrm{wd}'(v,w),w) \notin L_v$, then w.h.p. there exist $s_0,\dots,s_{j_0} = s'_w \in S$ such that*

- $\mathrm{wd}'(w,s'_w) \in (2+\mathcal{O}(\varepsilon))\,\mathrm{wd}_S(v,w)$ *and*
- $\mathrm{wd}'_S(v,s_0) + \sum_{j=1}^{j_0} \mathrm{wd}'_S(s_{j-1},s_j) \in (3+\mathcal{O}(\varepsilon))\,\mathrm{wd}(v,w)$.

Proof. By Statements 1 and 4 of Lemma 4.2, $(\mathrm{wd}'(v,w),w) \notin L_v$ means that $\mathrm{wd}'(v,s'_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,w)$ w.h.p. Hence,

$$\mathrm{wd}(v,s_v) \leq \mathrm{wd}(v,s'_v) \leq \mathrm{wd}'(v,s'_v) \leq (1+\varepsilon)\,\mathrm{wd}(v,w).$$

By the triangle inequality, it follows that

$$\mathrm{wd}(w,s_w) \leq \mathrm{wd}(w,s_v) \leq \mathrm{wd}(v,w) + \mathrm{wd}(v,s_v) \leq (2+\varepsilon)\,\mathrm{wd}(v,w).$$

w.h.p. Applying the second statement of Lemma 4.2 to w and s_w, we obtain that

$$\mathrm{wd}(w,s'_w) \leq \mathrm{wd}'(w,s'_w) \leq \mathrm{wd}'(w,s_w) \leq (1+\varepsilon)\,\mathrm{wd}(w,s_w) \leq 2(1+\varepsilon)^2\,\mathrm{wd}(v,w),$$

and, from the first statement,

$$\mathrm{wd}'(w,s'_w) \leq (1+\varepsilon)\,\mathrm{wd}(w,s'_w) \leq 2(1+\varepsilon)^3\,\mathrm{wd}(v,w),$$

i.e., the first part of the lemma's claim holds (recall that $\varepsilon \in \mathcal{O}(1)$). Moreover, it follows that

$$\mathrm{wd}(v, s'_w) \leq \mathrm{wd}(v, w) + \mathrm{wd}(w, s'_w) \leq 3(1+\varepsilon)^2 \,\mathrm{wd}(v, w).$$

Consider a shortest path from v to s'_w, and denote by $s_0, \ldots, s_{j_0} \in S$ the sampled nodes that are encountered when traversing it from v to s'_w; in particular, $s_{j_0} = s'_w$. By the same calculation as for Lemma 4.2, w.h.p. any two consecutive sampled nodes are no more than h hops apart. As the path is a shortest path from v to s'_w, the subpaths from s_{j-1} to s_j, $j \in \{1, \ldots, j_0\}$, and from v to s_0 are also shortest paths. Therefore, $h_{v,s_0} \leq h$ and, for each j, $h_{s_{j-1},s_j} \leq h$. We conclude that

$$\begin{aligned}\mathrm{wd}'_S(v, s_0) + \sum_{j=1}^{j_0} \mathrm{wd}'_S(s_{j-1}, s_j) &\leq (1+\varepsilon)\left(\mathrm{wd}(v, s_0) + \sum_{j=1}^{j_0} \mathrm{wd}(s_{j-1}, s_j)\right) \\ &= (1+\varepsilon)\,\mathrm{wd}(v, s'_v) \\ &\leq 3(1+\varepsilon)^3\,\mathrm{wd}(v, w),\end{aligned}$$

i.e., the second part of the claim of the lemma holds. □

Lemma 4.4. *Suppose we sample each node into S with independent probability p and solve $(1+\varepsilon)$-approximate (V, h, σ)-estimation with $h = c \log n/p$, where c is a sufficiently large constant. For $s \in S$, denote by T_s the tree induced by the routing paths from v to s for all $v \in V$ with $s'_v = s$. The depth of T_s is bounded by $\mathcal{O}(h \log n/\varepsilon)$, and each node participates in at most $\mathcal{O}(\log n)$ different trees.*

Proof. Recall that routing from v to s'_v is based on the routing tables $L_{v,i}$ determined by the unweighted source detection instances on G_i, $i \in \{0, \ldots, i_{\max}\}$. The induced shortest-path trees in G_i have depth at most $h' \in \mathcal{O}(h/\varepsilon)$, and they cannot overlap. By construction, the respective paths in G cannot have more hops. However, it is possible that when routing from v to s'_v, some node on the way knows of a shorter path to s'_v due to a source detection instance on G_j, $j \neq i$, and therefore "switches" to the shortest-path tree in G_j. Because $\mathrm{wd}_j(v, w) \geq \mathrm{wd}_i(v, w)$ for all $v, w \in V$ and $j \geq i$, we may however w.l.o.g. assume that the index i such that routing decisions are made according to $L_{v,i}$ is decreasing on each routing path from some node v to s'_v. Thus, the total hop count of the path is bounded by $\mathcal{O}(i_{\max} h') \subseteq \mathcal{O}(h \log n/\varepsilon)$. Consequently, the depth of each T_s is bounded by this value.

Concerning the number of trees a node may participate in, observe that if some node v decides that the next routing hop to s'_v is its neighbor u, it does so because s'_v minimizes the hop distance from v to s'_v in G_i, according to its list $L_{v,i}$. As there are $i_{\max} + 1 \in \mathcal{O}(\log n)$ different lists $L_{v,i}$, this is also a bound on the number of different trees v may participate in. □

We summarize with the following theorem.

Theorem 4.5. *For any $k \in \mathbb{N}$, routing table construction with stretch $6k - 1 + o(1)$ and labels of size $\mathcal{O}(\log n)$ can be solved in $\tilde{\mathcal{O}}(n^{1/2+1/(4k)} + D)$ rounds.*

Proof. We follow the approach outlined above, sampling nodes into S with probability $p = n^{-1/2-1/(4k)}$. By Chernoff's bound, this implies that $|S| \in \Theta(n^{1/2-1/(4k)})$ w.h.p. Using Corollary 3.5, we solve $(1+\varepsilon)$-approximate (V, h, σ)-estimation with $h = \sigma = c \log n/p$, for $c \in \mathcal{O}(1)$ sufficiently large, and, say, $\varepsilon = 1/\log n$. This takes $\tilde{\mathcal{O}}(1/p) = \tilde{\mathcal{O}}(n^{1/2-1/(4k)} + D)$ rounds and enables for each $v \in V$ to route to all nodes $w \in V$ with $(\mathrm{wd}'(v, w), w) \in L_v$ along a path of weight at most $\mathrm{wd}'(v, w)$. By Lemma 4.2, this enables for each $v, w \in V$ with $(\mathrm{wd}'(v, w), w) \leq (\mathrm{wd}'(v, s'_v), s'_v)$ to determine that this condition is satisfied and route from v to w with stretch $(1+\varepsilon)$.

To handle the possibility that $(\mathrm{wd}'(v,w),w) \notin L_v$, we call upon Corollary 3.5 once more. This time we solve $(1+\varepsilon)$-approximate $(S,h,|S|)$-detection; denote by wd'_S the corresponding distance function. Since $|S| \in \mathcal{O}(n^{1/2-1/(4k)})$ w.h.p., this requires $\tilde{\mathcal{O}}(n^{-1/2-1/(4k)})$ rounds w.h.p. We apply Lemma 4.3, showing that there are $s_0,\dots,s_{j_0} = s'_w \in S$ so that $\mathrm{wd}'(s'_w,w) \in (2+\mathcal{O}(\varepsilon))\,\mathrm{wd}(v,w)$ and $\mathrm{wd}'_S(v,s_0) + \sum_{j=1}^{j_0} \mathrm{wd}'_S(s_{j-1},s_j) \in (3+\mathcal{O}(\varepsilon))\,\mathrm{wd}(v,w)$. If we can route from v to s'_w incurring an additional stretch factor of $2k-1$ and from s'_w to w over a path of weight $\mathrm{wd}'(s'_w,w)$, the total stretch will be

$$(2+\mathcal{O}(\varepsilon)) + (2k-1)(3+\mathcal{O}(\varepsilon)) \in 6k-1+\mathcal{O}(\varepsilon) \subset 6k-1+o(1),$$

i.e., the routing scheme satisfies the claimed stretch bound.

Concerning routing from s'_w to w, we construct labels for tree routing using the algorithm from [20] that terminates in $\tilde{\mathcal{O}}(h)$ rounds in trees of depth h. By Lemma 4.4, this can be done in $\tilde{\mathcal{O}}(h) = \tilde{\mathcal{O}}(n^{-1/2-1/(4k)})$ rounds, where we simulate one round on each of the trees using $\mathcal{O}(\log n)$ rounds, one for each tree single node may participate in. The computed label of size $(1+o(1))\log n$ is added to the label of w, permitting to route from s'_w to w over a path of weight at most $\mathrm{wd}'(w,s'_w)$.

To route from v to s'_w, consider the graph on node set S with edge set $\{\{s,t\} \mid \mathrm{wd}'_S(s,t) < \infty\}$, where the edge weights are given by wd'_S. For this graph, each node $s \in S$ knows its incident edges and their weights. Using the simulation of the Baswana-Sen algorithm [3] given in [15], we can construct and make known to all nodes a $2k-1$ spanner[5] of this graph in

$$\tilde{\mathcal{O}}\left(|S|^{1+1/k}+D\right) = \tilde{\mathcal{O}}\left(n^{(1/2-1/(4k))(1+1/k)}+D\right) \subset \tilde{\mathcal{O}}\left(n^{1/2+1/(4k)}+D\right)$$

rounds. Using this knowledge, the fact that v knows $\mathrm{wd}'_S(v,s_0)$, and the routing tables from the second application of Corollary 3.5, w.h.p. we can route with the desired stretch from v to s'_w based on the identifier of s'_w, which we add to the label of v. This completes the proof of the stretch bound. Checking the individual bounds we picked up along the way, we see that the label size is $\mathcal{O}(\log n)$ and the running time is $\tilde{\mathcal{O}}(n^{1/2+1/(4k)}+D)$ w.h.p. □

4.3 Compact Routing on Graphs of Small Diameter

We now turn to the question of how to minimize the routing table size when computing routing tables distributedly. The idea in the algorithm is to construct an (approximate) Thorup-Zwick routing hierarchy [20]. We remark that compared to the original construction, we lose a factor of roughly 2 in stretch. This comes from the fact that in the centralized setting, one assumes access to the table of both nodes, i.e., the origin and destination of the routing or distance query. This is equivalent to identifying tables and lables, resulting in lable size $\tilde{\Theta}(n^{1/k})$. We consider this inexpedient for distributed systems and hence focus on obtaining small lables.

Our approach is efficient if D is small. Using exact distances, the construction would look as follows.

1. For each node $v \in V$, choose its *level* by an independent geometric distribution, i.e., the probability to have level at least $l \in \{0,\dots,k-1\}$ equals $p_l := n^{-l/k}$. Denote the set of nodes of level at least l by S_l; in particular, $S_0 = V$.
2. For each node v and each level $l \in \{1,\dots,k-1\}$, determine the closest node $s_l(v)$ to v and the set $S_{l-1}(v)$ of all nodes in S_{l-1} closer to v than $s_l(v)$ (ties broken by node identifiers); for notational convenience, let $s_0(v) := v$ and $S_{k-1}(v) := S_{k-1}$.
3. Determine tables and labels for routing and distance approximation (i) from v to all nodes in $S_l(v)$, for all $l \in \{0,\dots,k-1\}$, and (ii) from $s_l(v)$ to v, where $l \in \{1,\dots,k-1\}$. To determine the final label of v, concatentate its individual labels and the labels for routing from $s_l(v)$, $l \in \{1,\dots,k-1\}$.

[5] I.e., a subgraph in which distances increase by at most a factor $2k-1$.

In our implementation, we replace exact distances by $(1+\varepsilon)$-approximate distances for sufficiently small ε. Henceforth, we assume that the sets $S'_l(v)$ and nodes $s'_l(v)$ are defined as above, but with respect to wd'_l, the distance function corresponding to the instance of partial distance estimation we solve for level $l \in \{0,\dots,k-1\}$. Let us first examine the effect of the inaccurate distances on the stretch. This is done by a repeated application of the argument of Lemma 4.3.

Lemma 4.6. *For $l \in \{1,\dots,k-1\}$, denote by wd'_l the distance function corresponding to a $(1+\varepsilon)$-approximate solution to (S_l, h_l, σ_l)-estimation, where $h_l = \sigma_l = c\log n/p_l$ for a sufficiently large constant c. Suppose $v, w \in V$ and $\ell \in \{0,\dots,k-1\}$ is minimal so that $s'_\ell(w) \in S'_\ell(v)$. Then, w.h.p.,*

$$\mathrm{wd}(v, s'_\ell(w)) + \mathrm{wd}(s'_\ell(w), w) \le (1+\varepsilon)^{4\ell}(4\ell+1)\,\mathrm{wd}(v,w).$$

Proof. For $0 \le l \le \ell$, we prove by induction on l that

$$\mathrm{wd}(w, s'_l(w)) \le (1+\varepsilon)^{2l} 2l\,\mathrm{wd}(v,w)$$

and that

$$\mathrm{wd}(v, s'_l(w)) \le (1+\varepsilon)^{2l}(2l+1)\,\mathrm{wd}(v,w)$$

w.h.p. For $l=0$, trivially $\mathrm{wd}(v, s'_0(w)) = \mathrm{wd}(v,w)$ and $\mathrm{wd}(w, s'_0(w)) = \mathrm{wd}(w,w) = 0$. For the step from l to $l+1$, we make the intermediate claim that

$$\mathrm{wd}(v, s'_{l+1}(v)) \le (1+\varepsilon)^{2l+1}(2l+1)\,\mathrm{wd}(v,w).$$

If $\mathrm{wd}'_{l+1}(v, s'_l(w)) \le (1+\varepsilon)\,\mathrm{wd}(v, s'_l(w))$, then

$$\mathrm{wd}(v, s'_{l+1}(v)) \le \mathrm{wd}'_{l+1}(v, s'_{l+1}(v)) \le \mathrm{wd}'_{l+1}(v, s'_l(w)) \le (1+\varepsilon)\,\mathrm{wd}(v, s'_l(w))$$

w.h.p., where the second last step exploits that $s'_l(w) \notin S'_l(v)$ by the definition of $\ell > l$, but $s'_{l+1}(v) \in S'_l(v)$ w.h.p. Otherwise, Statements 1 and 4 of Lemma 4.2 yield that $\mathrm{wd}'_{l+1}(v, s'_{l+1}(v)) \le (1+\varepsilon)\,\mathrm{wd}(v, s'_l(w))$, resulting in the same bound on $\mathrm{wd}(v, s'_{l+1}(v))$. Either way, applying the induction hypothesis shows the claim.

We now can apply the triangle inequality to see that

$$\mathrm{wd}(w, s'_{l+1}(v)) \le \mathrm{wd}(v,w) + \mathrm{wd}(v, s'_{l+1}(v)) \le (1+\varepsilon)^{2l+1}(2(l+1))\,\mathrm{wd}(v,w).$$

If $\mathrm{wd}'_{l+1}(w, s'_{l+1}(v)) \le (1+\varepsilon)\,\mathrm{wd}(w, s'_{l+1}(v))$, then

$$\mathrm{wd}(w, s'_{l+1}(w)) \le \mathrm{wd}'_{l+1}(w, s'_{l+1}(w)) \le \mathrm{wd}'_{l+1}(w, s'_{l+1}(v)) \le (1+\varepsilon)\,\mathrm{wd}(w, s'_{l+1}(v)).$$

Otherwise, Statements 1 and 4 of Lemma 4.2 imply that $\mathrm{wd}'_{l+1}(w, s'_{l+1}(w)) \le (1+\varepsilon)\,\mathrm{wd}(w, s'_{l+1}(v))$ w.h.p., resulting in the same bound on $\mathrm{wd}(w, s'_{l+1}(w))$. Therefore, in both cases,

$$\mathrm{wd}(w, s'_{l+1}(w)) \le (1+\varepsilon)^{2(l+1)} 2(l+1)\,\mathrm{wd}(v,w),$$

i.e., the first part of the hypothesis is shown for index $l+1$. The second part readily follows by applying the triangle inequality once more, yielding

$$\mathrm{wd}(v, s'_{l+1}(w)) \le \mathrm{wd}(v,w) + \mathrm{wd}(w, s'_{l+1}(w)) \le (1+\varepsilon)^{2(l+1)}(2(l+1)+1)\,\mathrm{wd}(v,w)$$

w.h.p. This concludes the induction. Evaluating both statements of the induction hypothesis for index $l = \ell$ completes the proof. □

Lemma 4.6 shows that for $\varepsilon \in o(1/k)$, routing from v to w via $s'_\ell(w) \in S'_\ell(v)$ for minimal ℓ achieves stretch $4k-3+o(1)$. It remains to construct the hierarchy efficiently. We start with a general algorithm.

Lemma 4.7. *For each level $l \in \{0,\dots,k-1\}$, we can determine w.h.p. for all nodes v the set $S'_l(v)$ and the respective distance and routing information in $\tilde{\mathcal{O}}(\varepsilon^{-2}n^{(l+1)/k})$ rounds, where the tables have size $\mathcal{O}(n^{1/k}\log^2 n)$. Within this time, we can also determine labels of size $(1+o(1))\log n$ and tables of size $\mathcal{O}(\log^2 n)$ at each node for routing from $s'_l(v)$ to v.*

Proof. For a sufficiently large constant c, we perform $(1+\varepsilon)$-approximate (S_l, h_{l+1}, σ)-estimation with $h_{l+1} = cn^{(l+1)/k}\log n$ and $\sigma = cn^{1/k}\log n$. For $l < k-1$, the probability that $(\mathrm{wd}'(v, s'_{l+1}(v)), s'_{l+1}(v))$ has index $i \geq \sigma$ if we order $\{(\mathrm{wd}'(v,s),s) \mid s \in S_l\}$ ascendingly is $(1-p_l/p_{l+1})^\sigma \in n^{-\Omega(c)}$. The probability that $(\mathrm{wd}'(v, s'_{l+1}(v)), s'_{l+1}(v))$ has index $j \geq h_{l+1}$ if we order $\{(\mathrm{wd}'(v,w),w) \mid w \in V\}$ ascendingly is $(1-1/p_{l+1})^{h_{l+1}} \in n^{-\Omega(c)}$. By appending a bit to messages indicating whether $s \in S_l$ is also in S_{l+1}, we can thus use Corollary 3.5 to show that, w.h.p., we obtain suitable tables for routing from $v \in V$ to $S_l(v)$ and $s_{l+1}(v)$ within the stated time bound. If $l = k-1$, we have that $h_{l+1} > n$ and $|S_l| = |S_{k-1}| \leq \sigma$ w.h.p.; in this case, Corollary 3.5 shows that the construction can be performed as well.

Regarding the second part of the statement, observe that analogously to Lemma 4.4, the routing trees rooted at each node $s_{l+1} \in S_{l+1}$ have depth $\mathcal{O}(h_{l+1}/\varepsilon)$ and each node participates in at most $\mathcal{O}(\log n)$ of them. Thus, we can apply the construction from [20] to obtain labels (and tables) of size $(1+o(1))\log n$ for tree routing on each of the trees in $\tilde{\mathcal{O}}(h_{l+1}/\varepsilon) \subseteq \tilde{\mathcal{O}}(\varepsilon^{-2}n^{(l+1)/k})$ rounds. As each node participates in $\mathcal{O}(\log n)$ trees, the table size for this routing information is $\mathcal{O}(\log^2 n)$. □

Theorem 4.8. *Tables of size $\tilde{\mathcal{O}}(n^{1/k})$ and labels of size $\mathcal{O}(k\log n)$ for routing with stretch $4k-3$ can be computed in the* CONGEST *model in $\tilde{\mathcal{O}}(\mathit{SPD} + n^{1/k})$ rounds.*

Proof. We choose $\varepsilon \in \Theta(1/\log n)$. Then Lemma 4.6 shows that the stretch of the routing scheme will be $2(k-1)+2(k-1)+1 = 4k-3$. To construct the labels and tables, we use the approach of Lemma 4.7, but with $h := \mathrm{SPD}$. This is feasible, as by definition $h \geq h_{v,w}$ for all $v, w \in V$. As $\sigma \in \tilde{\Theta}(n^{1/k})$, and since we may assume w.l.o.g. that $k \in \mathcal{O}(\log n)$, the claimed running time bound follows. □

Unfortunately, Theorem 4.8 gives a good running time guarantee only when SPD is small. Worse, the strategy can be applied only if an upper bound on SPD is known (and the running time depends on that bound), unlike the algorithm of running time $\tilde{\mathcal{O}}(\mathrm{SPD}\cdot n^{1/k})$ from [5].[6] On the other hand, applying Lemma 4.7 to all levels (without modifying h) results in running time $\tilde{\mathcal{O}}(n)$. In the remainder of this section, we explain how to improve on Theorem 4.8 by "short-circuiting" the higher levels of the hierarchy. This approach yields better results when the hop diameter is small.

We now describe the construction. Let $l_0 < k-1$ be some level to be determined later. We will "truncate" the construction at level l_0 by constructing a *skeleton graph* as follows.

Definition 4.9 (l_0 skeleton graph)**.** *The* skeleton graph on level l_0 *is $G(l_0) = (S_{l_0}, E_{l_0}, \mathrm{wd})$, where $\{s,t\} \in E_{l_0}$ if and only if $h_{s,t} \leq cn^{l_0/k}\log n$ for a sufficiently large constant c. We denote $h_{l_0} := cn^{l_0/k}\log n$.*

The l_0 skeleton graph preserves the original skeleton distances, as the following lemma states.

Lemma 4.10. *For any $\varepsilon > 0$, $h, \sigma \in \mathbb{N}$, and $S \subseteq S_{l_0}$, denote by $\mathrm{wd}_{S_{l_0}}$ the distance function resulting from solving $(1+\varepsilon)$-approximate $(S_{l_0}, h_{l_0}, |S_{l_0}|)$-estimation and by wd_s the distance function resulting from*

[6] Their algorithm only handles distance queries and assumes that also the table of the destination can be accessed (i.e., the lables are identical to the tables). Both assumptions can be removed to achieve the same properties as our solution within $\tilde{\mathcal{O}}(\mathrm{SPD}\cdot n^{1/k})$ rounds.

solving $(1+\varepsilon)$-approximate $(S, ch\log n, \sigma)$-estimation on $G(l_0)$, where c is a sufficiently large constant. Then, w.h.p.,

$$\mathrm{wd}'(v,s) := \min_{t\in S_{l_0}}\{\mathrm{wd}'_{S_{l_0}}(v,t) + \mathrm{wd}'_S(t,s)\}$$

is a suitable distance function for $(1+\varepsilon)$-approximate $(S, h\cdot h_{l_0}, \sigma)$-estimation on G.

Proof. By the triangle inequality, for any $v\in V$, $t\in S_{l_0}$, and $s\in S$,

$$\mathrm{wd}(v,s) \le \mathrm{wd}(v,t) + \mathrm{wd}(t,s) \le \mathrm{wd}'_{S_{l_0}}(v,t) + \mathrm{wd}'_S(t,s).$$

Now suppose $h_{v,s} \le h\cdot h_{l_0}$ for some $v\in V$ and $s\in S$. The expected number of nodes in S_{l_0} on a shortest path from v to s of $h_{v,s}$ hops is $p_{l_0}h_{v,s} \in \mathcal{O}(h\log n)$. By Chernoff's bound, this number is smaller than $ch\log n$ w.h.p., as c is sufficiently large. Another application of Chernoff's bound shows that the maximum hop distance between nodes from S_{l_0} on the path is bounded by h_{l_0} w.h.p.

Denoting by $t_{v,s}\in S_{l_0}$ the first sampled node on the path, the above shows that the following properties hold w.h.p.

- $\mathrm{wd}(v,s) = \mathrm{wd}(v,t_{v,s}) + \mathrm{wd}(t_{v,s},s)$,
- $\mathrm{wd}_{G(l_0)}(t_{v,s},s) = \mathrm{wd}(t_{v,s},s)$, where $\mathrm{wd}_{G(l_0)}$ denotes the weighted distance in $G(l_0)$,
- $\mathrm{wd}'_{S_{l_0}}(v,t_{v,s}) \le (1+\varepsilon)\,\mathrm{wd}(v,t_{v,s})$, and
- $\mathrm{wd}'_S(t_{v,s},s) \le (1+\varepsilon)\,\mathrm{wd}_{G(l_0)}(t_{v,s},s)$.

Overall, this yields

$$\begin{aligned}
\mathrm{wd}'(v,s) &= \min_{t\in S_{l_0}}\{\mathrm{wd}'_{S_{l_0}}(v,t) + \mathrm{wd}'_S(t,s)\}\\
&\le \mathrm{wd}'_{S_{l_0}}(v,t_{v,s}) + \mathrm{wd}'_S(t_{v,s},s)\\
&\le (1+\varepsilon)(\mathrm{wd}(v,t_{v,s}) + \mathrm{wd}_{G(l_0)}(t_{v,s},s))\\
&= (1+\varepsilon)(\mathrm{wd}(v,t_{v,s}) + \mathrm{wd}(t_{v,s},s))\\
&= (1+\varepsilon)\,\mathrm{wd}(v,s).
\end{aligned}$$

□

Corollary 4.11. *If in the construction from Lemma 4.10 we replace $G(l_0)$ by the graph $\tilde{G}(l_0)$ constructed by solving $(1+\varepsilon)$-approximate $(S_{l_0}, h_{l_0}, |S_{l_0}|)$-estimation and assigning weight $\mathrm{wd}'_{S_{l_0}}(s,t)$ to edge $\{s,t\}$, the resulting function wd' is a suitable distance function for $(1+\varepsilon)^2$-approximate $(S, h\cdot h_{l_0}, \sigma)$-estimation.*

Proof. By definition, for all edges $\{s,t\}\in E_{l_0}$, we have that $\mathrm{wd}'_{S_{l_0}}(s,t) \le (1+\varepsilon)\,\mathrm{wd}(s,t)$. Also, clearly $\mathrm{wd}'_{S_{l_0}}(s,t) \ge \mathrm{wd}(s,t)$ for all $s,t\in S_{l_0}$. Therefore, we can reason analogously to Lemma 4.10, except for incurring another factor of $1+\varepsilon$ in stretch. □

Next, we consider the simulation of the truncated levels in the hierarchy.

Lemma 4.12. *For any integer $l_0 \ge k/2+1$, we can construct level $l\ge l_0$ of the routing hierarchy in $\tilde{\mathcal{O}}(\varepsilon^{-2}(n^{l_0/k} + n^{(k-l_0)/k}D))$ rounds w.h.p., where the tables and labels are of size $\mathcal{O}(n^{1/k})$ and $\mathcal{O}(\log n)$, respectively.*

Proof. Recall that $\varepsilon\in\mathcal{O}(1)$. We choose $\varepsilon'\in\Theta(\varepsilon)$ such that $(1+\varepsilon')^2 = (1+\varepsilon)$. We solve $(1+\varepsilon')$-approximate $(S_{l_0}, h_{l_0}, |S_{l_0}|)$-estimation using Corollary 3.5, w.h.p. requiring

$$\tilde{\mathcal{O}}(\varepsilon^{-2}(h_{l_0}+|S_{l_0}|)+D) = \tilde{\mathcal{O}}\left(\varepsilon^{-2}\left(n^{l_0/k} + n^{(k-l_0)/k}\right) + D\right) \subseteq \tilde{\mathcal{O}}\left(\varepsilon^{-2}n^{l_0/k} + D\right)$$

rounds. Our goal is to apply Corollary 4.11. To this end, we will simulate, for $h = h_{l+1}/h_{l_0}$ and sufficiently a sufficiently large constant c, $(1+\varepsilon')$-approximate $(S_l, c\,h\log n, c\,n^{1/k}\log n)$-estimation on $\tilde{G}(l_0)$, in a

way such that *all* nodes will learn the output of *all* nodes in S_{l_0}. As in Lemma 4.7, a bit indicating whether a source is in S_{l+1} is added to messages if $l < k-1$.

Before we explain how to do this, let us show how this permits to construct level l of the routing hierarchy. From the collected information, w.h.p. nodes can locally compute the distance function wd' from Corollary 4.11 for the σ closest nodes in S_l w.r.t. wd' and, analogously to Lemma 4.7, derive their table for routing from v to S'_l and $s'_l(v)$.

To enable tree routing from $s'_l(v)$ to v, split the tree rooted at $s'_l(v)$ into the unique maximal subtrees rooted at $s \in S_{l_0}$ that contain no internal nodes from S_{l_0} (i.e., all such nodes are either the root or leaves). By Lemma 4.4, these subtrees have depth at most $\tilde{\mathcal{O}}(h_{l_0})$. We use separate labeling schemes for the (globally known) tree on $\tilde{G}(l_0)$ that describes the connections between nodes in S_{l_0} in the routing tree rooted at $s'_l(v)$ and the subtrees rooted at each $s \in S_{l_0}$. The former can be computed locally. The latter can be labeled in time $\tilde{\mathcal{O}}(\varepsilon^{-2}h_{l_0})$, provided that each node participates in $\tilde{\mathcal{O}}(1)$ different trees only. Analogously to Lemma 4.4, this holds true because each routing decision must correspong to one of the $\mathcal{O}(\log n)$ top entries of the routing tables (either for routing in G to some node in S_{l_0} or in $\tilde{G}(l_0)$). This approach requires each node in the tree to store two labels of size $(1+o(1))\log n$. Routing can now be executed by determining the next node from S_{l_0} to visit on the path from $s'_l(v)$ to v (if there still is one) and then use the label for the current subtree to find the next routing hop.

It remains to discuss how to solve $(1+\varepsilon')$-approximate $(S_l, h, c\,n^{1/k}\log n)$-estimation on $\tilde{G}(l_0)$ quickly. Recall that each node in S_{l_0} knows its neighbors and the weights of incident edges from the solution of $(1+\varepsilon')$-approximate $(S_{l_0}, h_{l_0}, |S_{l_0}|)$-estimation computed earlier. We simulate the algorithm given by Corollary 3.5, exploiting that each node broadcasts in only $\tilde{\mathcal{O}}(n^{2/k})$ rounds in total. For each simulated round $i \in \{1,\dots,h+\sigma\}$, we pipeline the communication over a BFS tree, which takes $\mathcal{O}(M_i + D)$ rounds in G, where M_i is the number of nodes in $\tilde{G}(l_0)$ that broadcast in simulated round i; this time bound includes $\mathcal{O}(D)$ rounds for global synchronization of when the next simulated round starts. Therefore, the total number of communication rounds in G is

$$\sum_{i=1}^{h+\sigma}\mathcal{O}(M_i+D) \subseteq \tilde{\mathcal{O}}(\sigma^2|S_{l_0}|+(h+\sigma)D) \subseteq \tilde{\mathcal{O}}(n^{2/k}\cdot n^{(k-l_0)/k}+n^{(l-l_0+1)/k}D)) \subseteq \tilde{\mathcal{O}}(n^{l_0/k}+n^{(k-l_0)/k}D))$$

w.h.p., as $|S_{l_0}| \in \tilde{\mathcal{O}}(n^{(k-l_0)/k})$ w.h.p. The bounds on table and lable size follow from Lemma 4.7 and the above discussion of the tree labeling scheme. □

We can now put all the pieces together to obtain the following result.

Theorem 4.13. *Suppose we are given $k \in \mathbb{N}$ and some integer $k/2+1 \le l_0 \le k$. Then tables of size $\tilde{\mathcal{O}}(n^{1/k})$ and labels of size $\mathcal{O}(k\log n)$ enabling routing and distance approximation with stretch $4k-3+o(1)$ can be constructed in $\tilde{\mathcal{O}}(n^{l_0/k}+n^{(k-l_0)/k}D)$ rounds w.h.p.*

Proof. Fix $\varepsilon := 1/\log^2 n$. W.l.o.g., $k \in \mathcal{O}(\log n)$. We construct the first l_0 levels of the hierarchy using Lemma 4.7, and the remaining levels using Lemma 4.12. This yields the stated running time, by Lemma 4.6 and Corollary 4.11 stretch $(1+\varepsilon)^{\mathcal{O}(k)}(4(k-1)+1) \in 4k-3+o(1)$, table size $\tilde{\mathcal{O}}(kn^{1/k}) = \tilde{\mathcal{O}}(n^{1/k})$, and labels of size $\mathcal{O}(k\log n)$. □

We can now choose an appropriate value for l_0 depending on D. If the running time is worse than about $n^{2/3}$, we handle the higher levels simply by making $\tilde{G}(l_0)$ known to all nodes and solving locally.

Corollary 4.14. *For any integer $k > 1$, tables of size $\tilde{\mathcal{O}}(n^{1/k})$ and labels of size $\mathcal{O}(k\log n)$ enabling routing and distance approximation with stretch $4k-3+o(1)$ can be constructed in*

$$\tilde{\mathcal{O}}\left(\min\{(Dn)^{1/2}\cdot n^{1/k}, n^{2/3+2/(3k)}\}+D\right)$$

rounds w.h.p.

Proof. If $k = 2$, the minimum is attained for the second term. If $k \geq 3$, choose l_0 as the closest integer to $k(\log D/\log n + 1)/2$, however, at least $k/2 + 1$ and at most $k - 1$. By Theorem 4.13, we can compute compact routing tables within

$$\tilde{\mathcal{O}}\left(n^{l_0/k} + n^{(k-l_0)/k}D\right) \subseteq \tilde{\mathcal{O}}\left((Dn)^{1/2}n^{1/k} + D \cdot n^{1/k}\right),$$

where the first summand covers the case that $l_0 \neq k - 1$ and the second one the possibility that $l_0 = k - 1$, implying that $l_0 \geq k - 3/2$ and hence $Dn^{1/k} \geq n^{1-1/(2k)} > n^{l_0/k}$. Note that in the latter case, we have that $Dn^{1/k} > n^{5/6}$, i.e., the term is irrelevant for the minimum in the running time bound of the corollary unless $k = 3$.

Alternatively, we may handle levels $1, \ldots, l_0 - 1$ using Lemma 4.7, determine $\tilde{G}(l_0)$, broadcast all its edges over a BFS tree, and compute the table construction locally. This takes

$$\tilde{\mathcal{O}}(n^{l_0/k} + n^{2(k-l_0)} + D)$$

rounds w.h.p., as $\tilde{G}(l_0)$ has at most $|S_{l_0}|^2$ edges and $|S_{l_0}| \in \mathcal{O}(n^{(k-l_0)/k})$ w.h.p. For the optimal choice of l_0, this results in running time

$$\tilde{\mathcal{O}}(n^{2/3+2/(3k)} + D).$$

For the special case of $k = 3$, note that $l_0 = 2$ in fact yields running time $\tilde{\mathcal{O}}(n^{2/3} + D) \subset \tilde{\mathcal{O}}((nD)^{1/2}n^{1/3})$. Hence we may drop the additive term of $n^{1/k}D$ from the first bound when taking the minimum. □

References

[1] I. Abraham, C. Gavoille, D. Malkhi, N. Nisan, and M. Thorup. Compact name-independent routing with minimum stretch. *ACM Trans. Algorithms*, 4(3):37:1–37:12, 2008.

[2] B. Awerbuch, A. Bar-Noy, N. Linial, and D. Peleg. Compact distributed data structures for adaptive network routing. In *Proc. 21st ACM Symp. on Theory of Computing*, pages 230–240, May 1989.

[3] S. Baswana and S. Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Structures and Algorithms*, 30(4):532–563, 2007.

[4] R. E. Bellman. On a routing problem. *Quart. Appl. Math.*, 16:87–90, 1958.

[5] A. Das Sarma, M. Dinitz, and G. Pandurangan. Efficient computation of distance sketches in distributed networks. In *Proc. 24th ACM Symp. on Parallelism in Algorithms and Architectures*, 2012.

[6] A. Das Sarma, S. Holzer, L. Kor, A. Korman, D. Nanongkai, G. Pandurangan, D. Peleg, and R. Wattenhofer. Distributed Verification and Hardness of Distributed Approximation. *SIAM Journal on Computing*, 41(5):1235–1265, 2012.

[7] L. R. Ford. Network flow theory. Technical Report P-923, The Rand Corp., 1956.

[8] S. Frischknecht, S. Holzer, and R. Wattenhofer. Networks cannot compute their diameter in sublinear time. In *Proc. 23rd ACM-SIAM Symp. on Discrete Algorithms*, pages 1150–1162, 2012.

[9] S. Holzer and N. Pinsker. Approximation of Distances and Shortest Paths in the Broadcast Congest Clique. *CoRR*, abs/1412.3445, 2014.

[10] C. Lenzen and D. Peleg. Efficient distributed source detection with limited bandwidth. In *Proc. 32nd ACM Symp. on Principles of Distributed Computing*, 2013.

[11] J. McQuillan, I. Richer, and E. Rosen. The new routing algorithm for the ARPANET. *IEEE Trans. Communication*, 28(5):711–719, May 1980.

[12] J. M. McQuillan and D. C. Walden. The ARPANET design decisions. *Networks*, 1, 1977.

[13] J. Moy. OSPF version 2, April 1998. Internet RFC 2328.

[14] D. Nanongkai. Distributed Approximation Algorithms for Weighted Shortest Paths. In *Proc. 46th Symposium on Theory of Computing (STOC)*, pages 565–573, 2014.

[15] B. Patt-Shamir and C. Lenzen. Fast Routing Table Construction Using Small Messages [Extended Abstract]. In *Proc. 45th Symposium on the Theory of Computing (STOC)*, 2013. Full version at http://arxiv.org/abs/1210.5774.

[16] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA, 2000.

[17] D. Peleg and A. A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989.

[18] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, 36(3):510–530, 1989.

[19] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.

[20] M. Thorup and U. Zwick. Compact routing schemes. In *Proc. 13th ACM Symp. on Parallel Algorithms and Architectures*, 2001.

[21] U. Zwick. Exact and approximate distances in graphs - a survey. In *Proc. 9th European Symp. on Algorithms*, pages 33–48, 2001.

[22] U. Zwick. All pairs shortest paths using bridging sets and rectangular matrix multiplication. *Journal of the ACM*, 49(3):289–317, 2002.